

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

SYSTEM AND METHOD FOR INQUIRY CACHING

Background of Invention

- [0001] Field of the invention. The invention relates generally to network systems and more particularly to systems and methods implemented in a network having a plurality of client devices that may submit commands to the same server device, wherein commands for relatively static data may be blocked until a previous command is serviced.
- [0002] Background. A computer network is comprised essentially of a group of devices that are coupled together by a communication medium. A network may be as simple as two devices that are wired to each other, or as complex as the Internet. The devices in a network can communicate with each other and provide services to each other (or, conversely, use each other's services).
- [0003] The devices may be of many different types. For example, some may be host devices (hosts) that initiate commands, while others may be devices which are designed to service the hosts (e.g., data storage devices). Networks may be configured in many different ways. For instance, referring to FIGURE 1, one particular configuration of a storage area network (SAN) is illustrated. It should be noted that a SAN may be configured in many different ways, and that the configuration illustrated in FIGURE 1 is exemplary.
- [0004] In the SAN of FIGURE 1, a plurality of hosts 11-13 are coupled to a sequential device 15 by a transmission medium 14. ("Sequential device," as used here, refers to a device that is sequentially accessed, such as a tape drive, rather than one which can be randomly accessed, such as a disk drive.) Sequential device 15 is configured to sequentially process requests for data which are received from hosts 11-13.

- [0005] Hosts 11-13 may transmit several different kinds of commands to sequential device 15. one type of command is an inquiry command. "Inquiry command," for the purposes of this disclosure, will refer to commands which request data that, for the most part, is static. This data, which will be referred to herein as "inquiry data,"typically relates to the device itself, rather than data which is designed to store or generate. For instance, inquiry data may include the serial number, manufacturer, configuration, version number, or similar data. It should be noted that, however, that inquiry data may also include data that changes relatively infrequently, rather than being strictly static.
- [0006] Inquiry commands may be submitted to device 15 for various reasons. For example, when a new host is booted, it checks to see what other devices are connected to the network. Hosts may also periodically use inquiry commands to obtain information regarding the availability of devices on the network. If a device responds to the command, the host will receive the responsive inquiry data and will be aware that the device is available. If the device does not respond, the host may assume either that the device is no longer connected to the network, or that the device is no longer functioning properly.
- [0007] One of the problems with handling inquiry commands in sequential devices is that these commands may be submitted to a sequential device just after the device has received a command that will take a substantial amount of time to process. Conventionally, no distinction is made between inquiry commands and other commands, so the inquiry command in this example will be handled in the same manner as any other command it will be serviced when the device becomes available. Until the previously received command is serviced, the inquiry command will be blocked. If the inquiry command is blocked for long enough, it will time out. As a result of having timed out, the host may assume that the device is no longer operational or no longer connected to the network, even though it is both connected and operational.
- [0008] It would therefore be desirable to provide a means for servicing inquiry commands in a sequential device while other commands are being processed by the device.

Summary of Invention

- [0009] One or more of the problems outlined above may be solved by the various embodiments of the invention. Broadly speaking, the invention comprises systems and methods for storing inquiry data for a device so that inquiry commands to the device can be serviced when the device is busy processing a previously received command.
- [0010] In one embodiment, a router contains a memory for storing static inquiry data corresponding to each of a plurality of sequential devices connected thereto. The first time a host sends an inquiry command to one of the sequential devices, the router routes the command to the target device and stores the data returned by the device. Subsequent inquiry commands are serviced by passing the command through to the target device if it is not busy, or returning inquiry data from the router's memory if the target device is busy processing another command. The router may be configured to determine whether the target device is busy in a number of ways, For example, it may track whether data has been returned in response to a command previously routed to the target device. The router may alternatively route all inquiry commands to the target device and, if the command is not serviced in a predetermined amount of time, the router may service the command from memory. Once the inquiry data has been stored in the router's memory, subsequent inquiry commands may be serviced either by always returning the data from the router's memory, or by returning the stored data only when the target device is busy and otherwise passing the command through to the target device so that it can service the request.
- [0011] Another embodiment comprises a method in which inquiry data is stored in a memory apart from the corresponding target device, an inquiry command is subsequently received, and the inquiry command is serviced by returning the corresponding inquiry data from the memory. The method may further comprise routing an initial inquiry command to the target device, regardless of whether the target device is busy, and continuing to do so until the target device returns responsive inquiry data, which is then stored in the memory. The method may include tracking whether data has been returned in response to a command previously routed to the target device to determine whether a target device is busy, or it may comprise routing all inquiry commands to the target device and servicing the commands from memory if they are not serviced in a predetermined amount of time. The method may comprise always servicing inquiry commands by returning the data from the router's

memory, or it may comprise returning the stored data only when the target device is busy.

[0012] Other embodiments may include computers or other data processing devices, computer readable media (such as floppy disks, CD-ROMs, DVD-ROMs, etc.,) storage devices, computer memories and the like which contain software, firmware or other programming which embodies the foregoing methods. Numerous other alternative embodiments are also possible.

Brief Description of Drawings

[0013] Other objects and advantages of the invention may become apparent upon reading the following detailed description and upon reference to the accompanying drawings.

[0014] FIGURE 1 is a diagram illustrating a plurality of host devices coupled to a sequential device in one embodiment.

[0015] FIGURE 2 is a diagram illustrating a system having a router configured to route the requests of several host devices to a sequential device.

[0016] FIGURE 3 is a diagram illustrating the routing of inquiry commands to a sequential device, the transmission of the responsive inquiry data and the storage of the inquiry data in a cache.

[0017] FIGURE 4 is a diagram illustrating the apparent configuration of a router and a sequential device with respect to a device that submits an inquiry command while the sequential device is busy processing a previous command.

[0018] FIGURE 5 is a flow diagram illustrating the operation of a router in accordance with one embodiment of the present invention.

[0019] FIGURE 6 is a flow diagram illustrating the operation of a router in accordance with an alternative embodiment of the present invention.

[0020] While the invention is subject to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and the accompanying detailed description. It should be understood, however, that the

drawings and detailed description are not intended to limit the invention to the particular embodiment which is described. This disclosure is instead intended to cover all modifications, equivalents and alternatives falling within the scope of the present invention as defined by the appended claims.

Detailed Description

- [0021] As noted above, host devices generally use the inquiry command to obtain information regarding the availability of devices on the network. If a host is booted up and a device on the network is busy, the device can't respond to the inquiry. A mechanism is therefore necessary to provide the responsive data, thereby keeping the host's inquiry command from timing out and keeping the host from assuming that the device is not available.
- [0022] Referring to FIGURE 2, an exemplary system in accordance with one embodiment of the invention is shown. In this embodiment, a plurality of hosts 21-23 are coupled to a router 25. The hosts are configured to communicate with the router over a Fibre Channel network 24. Router 25 is configured to convert data received from hosts 21-23 from the Fibre Channel protocol to the SCSI protocol. The data is then communicated over SCSI bus 26 to sequential device 27.
- [0023] When one of hosts 21-23 needs to send a command to sequential device 27, it transmits the command to router 25, which then forwards the command to sequential device 27. If sequential device 27 is not processing any other commands, it is available to process the newly received command. If sequential device 27 is processing a previous command, the subsequently received command cannot be processed until the previous command is completed.
- [0024] Router 25, however, is coupled to a cache 28. Cache 28 is designed to store inquiry data associated with sequential device 27. If sequential device 27 is busy processing a command when an inquiry command is received by router 25, the router is configured to examine the inquiry command, determine the data which is needed to respond to the command, and transmit the responsive data back to the host that generated the command.
- [0025] This assumes, of course, that the cache contains the data necessary to respond to

the inquiry command. When the router is booted, the cache contains no data. Consequently, it cannot provide any inquiry data in response to an inquiry command. If the router receives an inquiry command when the responsive data is not in the cache, and the sequential device itself is not available to provide responsive data, the result is the same as for prior art systems the inquiry command is blocked and, if the device does not become available to service the command soon enough, the command will time out. If the sequential device becomes available in time, it will service the command.

[0026] The first time the sequential device provides inquiry data to one of the hosts, the data is stored by the router in its cache. This is illustrated in FIGURE 3. Thereafter, the router can provide the data in response to inquiry commands if the sequential device is busy. The router can, in some embodiments, be configured to provide the data in response to inquiry commands if the sequential device is offline, broken, missing, or cannot respond for some other reason. In one embodiment, an indicator is set in the inquiry data when the device cannot respond for some reason other than being busy. If the inquiry data is returned with this indicator set, the requesting host knows that the inquiry data is the information that was requested, but the device is currently missing or offline. This configuration may be useful in systems that have been altered to allow a device to be repaired and later re-inserted into the system without having to reconfigure the hosts to look for a device change. A host operator can therefore simply avoid using the device for a short period and then use it again when it has been replaced.

[0027] Referring to FIGURE 4, a diagram illustrating the effect of the present system and method is shown. In this figure, two host devices, 33 and 34, are connected to router 31. Router 31 is, in turn, connected to sequential device 32. As depicted in the figure, host device 33 issues a command to sequential device 32. In this particular instance, there are no other commands pending with sequential device 32 (host device 34 has not yet issued a command). Router 31 therefore forwards the command to sequential device 32, which begins to process the command. (The command data path is indicated by arrow 35.)

[0028] It is assumed that the command issued by host device 33 is not an inquiry

command and it takes a certain amount of time for sequential device 32 to process. While this first command is being processed by sequential device 32, host device 34 issues a command. The command issued by host device 34 is an inquiry command. When this command is received by router 31, the router detects that sequential device 32, which is the target of the command, is busy. Router 31 therefore accesses its cache (not shown) and retrieves the inquiry data necessary to service the inquiry command. (This command data path is indicated by arrow 36.) To host device 34, it appears that its inquiry command was serviced by sequential device 32.

[0029] Referring to FIGURE 5, a flow diagram illustrating the operation of one embodiment of the present invention is shown. This figure corresponds to the functionality of a router as described above. First, a command is received from one of the host devices. The router must then determine whether or not the target sequential device is busy. If that target device is not busy, the command is forwarded to the target device. After that target device processes that command, it provides data responsive to the command. This data is transmitted back to the router. The router does two things with the data: it stores the data in its cache; and it forwards the data to the host device that originally requested it. It should be noted that, if the return to data is truly static, the router need only store the data the first time it is returned. It is possible to skip the step of storing the data following subsequent responses to the same inquiry command. It is contemplated, however, that for purposes of simplicity, it will be preferred to store the data each time it is returned.

[0030] If, when the router receives the command from the host device, the target device is busy, the router will respond to the request if possible. In order to provide a response, the router must have the data necessary to service the request stored in its cache. It therefore checks the cache to determine whether or not it has this data. If the data is stored in the cache, the data is read from the cache and then forwarded to the host device in response to the inquiry command. If the data is not stored in the cache, the command is held by the router until the target device becomes available. If the command does not time out before the target device becomes available, it will be forwarded to the target device. When the target device returns data responsive to the inquiry command, the data will be stored in the cache of the router and forwarded to the host device that requested it.

[0031] It should be pointed out that some provision should be made for determining when a target device has been disconnected from the router. If no mechanism is provided for this purpose, the router and host devices may simply assume that the target device is busy. In this case, the host devices may continue to make futile attempts to access the device. Possibly worse, the router may continue to provide inquiry data in response to commence from the host devices, thereby perpetuating the fallacy that the target device is still available. If, on the other hand, the device is identified as having been disconnected, the router and host devices can identify the target device as such. As a result, the host devices will not waste time attempting to access the target device, and the router will not erroneously provide inquiry data for a non-existent member of the network.

[0032] Referring to FIGURE 6, a flow diagram illustrating the operation of an alternative embodiment of the present invention shown. This diagram is, in many respects, similar to the diagram of FIGURE 5. In the diagram of FIGURE 6, however, the router's first action with respect to any received inquiry command is to determine whether or not data responsive to the command is stored in its cache. If the data is in the cache, it is read from the cache and returned to the host device without accessing the target device. If the data is not in the cache, the router must determine whether or not the target device is busy. If the target device is not busy, the inquiry command is forwarded to it. The data returned by the target device would then be forwarded to the requesting host device, as well as being stored in the router's cache for use in servicing future inquiry commands. If the target device is busy, the inquiry command is held until the target device becomes available. The inquiry command is then processed by the target device if it has not timed out.

[0033] The advantage of first checking the cache for the data is that the target device no longer has to deal with it. The command can be handled entirely by the router. The disadvantage of this embodiment is that it cannot be used when the inquiry data is slowly changing, as opposed to static. In the embodiment of FIGURE 5, slowly changing data would be updated whenever a host device issued an inquiry command and the target device was not busy (and therefore responded to the command). In the embodiment of FIGURE 5, the data which is originally stored in the cache in response to the inquiry command would be returned every time the inquiry command was

received, regardless of whether or not the data that would be provided by the target device itself had changed. Thus, the embodiment of FIGURE 6 would be sufficient if the inquiry data never changed, but not if the data were slowly changing.

[0034] It should be noted that the flow diagrams of FIGURES 5 and 6 do not explicitly include the step of determining whether a received command comprises an inquiry command or other command requesting static or near-static data. If the received command is a request for sequential data stored on the device, there is no need to process the command according to the flow diagrams, since the data will not be stored in the same manner as inquiry data. The determination that a received command is a request for sequential data will therefore cause the remainder of the flow diagram to be bypassed. If the received command is an inquiry command, the method of the flow diagram will proceed.

[0035] It should be noted that, although the embodiments described above are directed to systems in which a router is used to couple a set of Fibre Channel hosts to a SCSI device, the present systems and methods are not limited to use in Fibre Channel-to-SCSI systems, nor is it necessary that there be any protocol conversion at all. Other embodiments may use devices coupled between communication media having other protocols, or even between media having protocols which are identical to each other. Further, it is not necessary that the systems/methods be implemented in a router. Other embodiments may employ alternative types of devices to perform the functions of the routers described above (i.e., the storage and return of inquiry data). For example, this functionality may be incorporated into circuitry within the sequential device. Another embodiment may comprise a software application (and/or the media on which the application is stored) for controlling a router or other device to provide the desired functionality.

[0036] The benefits and advantages which may be provided by the present invention have been described above with regard to specific embodiments. These benefits and advantages, and any elements or limitations that may cause them to occur or to become more pronounced are not to be construed as a critical, required, or essential features of any or all of the claims. As used herein, the terms "comprises," "comprising," or any other variations thereof, are intended to be interpreted non-

exclusively, including the elements or limitations which follow those terms. Accordingly, a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to the claimed process, method, article, or apparatus.

[0037] While the present invention has been described with reference to particular embodiments, it should be understood that the embodiments are illustrative and that the scope of the invention is not limited to these embodiments. Many variations, modifications, additions and improvements to the embodiments described above are possible. For example, these variations may include computers or other data processing devices, computer readable media (such as floppy disks, CD-ROMs, DVD-ROMs, etc.,) storage devices, computer memories and the like which contain software, firmware or other programming which embodies the foregoing methods. It is contemplated that these variations, modifications, additions and improvements fall within the scope of the invention as detailed within the following claims.